arXiv:1702.07368v1 [q-bio.NC] 23 Feb 2017

Reconfiguring motor circuits for a joint manual and BCI task

Benjamin Lansdell[1+ *], Ivana Milovanovic[2+], Cooper Mellema[3], Eberhard E Fetz[3,4,5], Adrienne L Fairhall[3,4,5], Chet T Moritz[2,3,4,5],

1 Applied Mathematics, University of Washington, Seattle, WA, USA
2 Rehabilitation Medicine, University of Washington, Seattle, WA, USA
3 Physiology and Biophysics, University of Washington, Seattle, WA, USA
4 Center for Sensorimotor Neural Engineering, University of Washington, Seattle, WA, USA
5 UW Institute for Neural Engineering, University of Washington, Seattle, WA, USA

+ These authors contributed equally to this work.
* lansdell@uw.edu

Abstract

Designing brain-computer interfaces (BCIs) that can be used in conjunction with ongoing motor behavior requires an understanding of how neural activity co-opted for brain-control interacts with existing neural circuits. BCIs may be used to regain lost motor function after stroke or traumatic brain injury, for instance, provided that a neural control signal can be extracted from a single spared hemisphere. Such an approach, however, requires that neural activity controlling the unaffected limb is dissociated from activity controlling the BCI. In this study we investigate how primary motor cortex accomplishes simultaneous BCI control and motor control in a task that explicitly requires both activities to be driven from the same brain region (i.e. a 'dual control task'). Single unit activity was recorded from intracortical multi-electrode arrays while a non-human primate performed this dual-control BCI task. We observe broad changes in tuning of both units used to drive the BCI directly ('control units') and units that do not directly control the BCI ('non-control units'). Using transfer entropy, a measure of effective connectivity, we observe control-unit-specific dissociation from other units. Through an analysis of variance we find that the intrinsic variability of the control units has a significant effect on task proficiency. In particular, at least one control unit must have a high intrinsic variability to allow cursor control and target acquisition. Surprisingly, the degree of co-tuning or connectivity with other units did not affect dual-control task proficiency. Thus, motor cortical activity is flexible enough to perform novel BCI tasks that require active decoupling of natural associations to wrist motion.

Author summary

State-of-the-art brain-computer interface (BCI) control is achieved via intra-cortical recordings, in which individual motor cortical neurons provide signals to control a neural prosthetic. Previous single-unit BCI studies showed that primary motor cortex exhibits remarkable flexibility: selection of neurons for prosthetic control can be performed largely independently of their prior association with motor behavior. However, most such studies do not involve simultaneous ongoing motor output. In future applications, BCIs may need to be used simultaneously with other motor behavior, such as when only one hemisphere remains intact following stroke or traumatic brain injury. Here we investigate a dual-control task which explicitly requires ongoing wrist motion in conjunction with BCI controlled by the same region of cortex. By analyzing both the neurons that directly control the BCI and those that do not, we find that population-wide changes in tuning to wrist

movement support dual control. In addition, the task induces control-unit specific changes in effective connectivity. Only the control units' intrinsic variability predicts task proficiency, suggesting that the flexibility of motor cortical circuits can be leveraged for concurrent control of a BCI and natural movement.

Introduction

Broad application of brain-computer interfaces (BCIs) for control of neural prostheses requires understanding how brain circuits can simultaneously engage in competing tasks. The complexity of motor cortical neuron representations, however, creates uncertainty about how BCI control could be decoupled from ongoing motor control. Concurrent control of natural movement and BCI output may be useful in cases of brain injury. For example, when the motor cortex of one hemisphere is damaged by stroke or trauma, a BCI-controlled neural prosthetic may be driven by the intact hemisphere. In this case cortical circuits in the intact hemisphere may be tasked with both producing movement of the unaffected limb, while also restoring function to the affected, ipsilateral limb using a BCI. Even in the absence of injury, future industrial applications may seek to control supernumerary artificial limbs or external tools using cortex presumably already involved in normal behavior.

Current state-of-the-art BCI control is achieved via intra-cortical brain signals. Typical BCIs in humans and animals are derived from population coding based on real or imagined movement [1–4] . BCI studies observe that brain-control mappings which utilize activity observed during the natural motor repertoire are most effective [5, 6]. However utilizing natural mappings between population activity and movement may pose a challenge when BCI control is required to be independent of ongoing motor activity.

A possible approach, then, is to take advantage of biofeedback and conditioning paradigms that allow for volitional control of individual neurons. Studies demonstrate that monkeys can be trained to volitionally control individual neurons in such a fashion [7–11], with similar results also observed in humans [1, 3]. In these tasks, there is evidence that neural activity can be chosen independently of natural movement association [9–12]; furthermore, motor cortex cell activity can be operantly dissociated from activity of associated muscles [13], reflecting the remarkable adaptability of motor cortical activity. These findings suggest that next-generation BCI design may be able to better leverage the brain's inherent adaptability.

Most BCI studies to date focus on control of external devices alone. In these tasks, the subjects do not concurrently perform motor actions, thus freeing the activity of neurons or single units controlling the BCI to be modified as needed to perform the task. Indeed, in these cases, continued use of motor cortex cells to control a BCI resulted in a decrease in correlation to movement, suggesting decoupling of the cells from motor control [13, 14]. It is unknown whether and how such a decoupling occurs when natural motor control is engaged concurrently with BCI tasks.

While subpopulations of neurons encoding ipsilateral motion are a candidate control source for the BCI following stroke [15], a robust interface may require recording from the larger population of neurons involved in natural, contralateral control. When a single brain region must simultaneously control a BCI and ongoing limb movement, do neurons strongly associated with contralateral motion make poor control units due to the potential interference imposed by concurrent hand movement? Assuming that the cortical circuits sampled by the recording array have some causal role in motor control, it seems that at least some sampled neurons could present challenges for simultaneous manual and BCI control.

Previous studies of the activity of units used in control (control units) and the non-control units in the surrounding population (non-control units) reveal widespread changes in activity and tuning preferences in a variety of tasks and task perturbations [14, 16–21]. Hwang et al 2013 [5] observed changes in preferred direction of both control and non-control units when a BCI decoder's output is rotated, suggesting a global reconfiguration of the motor cortical network. In addition to these changes, other studies showed specific changes for the control units compared with those

not involved in the control [12,20]. How do these unit specific effects manifest in a BCI paradigm where control units may be additionally constrained by their role in ongoing movement?

In this study we directly address these questions using a 'dual-control' BCI task. In this task both BCI control and ongoing motor output are driven by the same cortical region, requiring control units to dissociate their activity from the circuits controlling the movement. The task requires that a monkey acquire a target in two dimensions, one dimension controlled by a BCI while the other is controlled using natural motor control of the wrist contralateral to the recorded hemisphere (Fig 1).

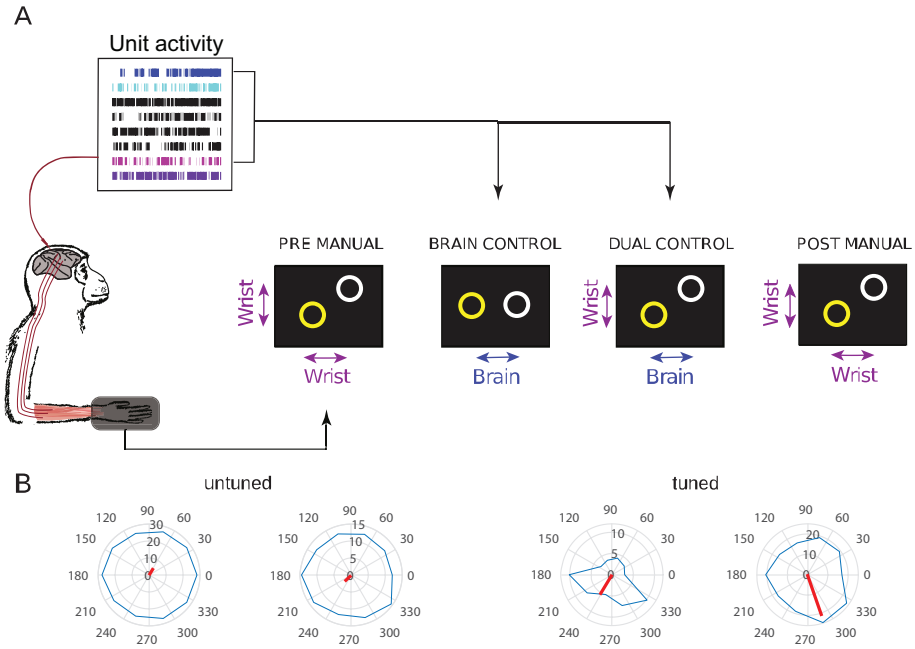

Fig 1. Dual control experimental setup. A: Dual control experimental setup. Isolated primary motor unit activity controls a brain-control axis, while contralateral wrist torque determines a manual-control axis. Monkey first performed a 2D manual wrist task. The monkey then performed a 1D brain control task, in which cursor velocity was determined by neural activity of two units. Then the monkey performed a dual control task, where one axis is determined by wrist torque and the other by neural activity. Finally, the monkey performed a second 2D manual task. B: Linear encoding models were used to identify preferred direction and modulation depth (tuning angle and tuning strength) of each unit. Two untuned units shown left have firing rate independent of wrist velocity direction, while two tuned units shown right respond preferentially to wrist velocity in particular directions. Radius indicates mean firing rate (Hz). Red line indicates tuning angle and strength.

The dual control task allows us to investigate how primary motor cortical neural activity can be decoupled from natural motion. A previous dual-control BCI study demonstrated that task performance was independent of control units' prior association to wrist motion [22]. Other concurrent-use BCIs demonstrate similar robustness, albeit often with some adaptation period required [23,24]. Here we extend these studies by recording activity of both control and non-control units during dual control. We examine whether dissociation from wrist motion occurs specifically for control units, and whether this dissociation is associated with changes in tuning and intrinsic variability at the population level. We also consider the *effective connectivity* of neurons in the network, which quantifies the extent to which firing of each neuron can be predicted from that of the others. Using these measures, we compare the motor cortical activity observed during the dual-control task to a standard brain-control task which requires the monkey to acquire targets only through BCI control.

Along with the clinical goals for BCI development, studying brain plasticity in the presence of

a BCI has the potential to elucidate the nature of motor network function [25]. While traditional approaches propose that motor cortical activity encodes motion kinetics – a representational view of cortical activity– contemporary theories consider motor cortical activity to be better described as a high dimensional interacting system in which is embedded a potentially lower-dimensional dynamical system whose features are related to muscle activities directly [26–29]. Here, our goal is to characterize as effectively as possible any changes in the population neural activity underlying a dual-control and non dual-control BCI. Thus, we examine both the properties of traditional representational features such as tuning curves, as well as dynamical features related to effective connectivity and low dimensional manifolds. This allows us to draw on a significant BCI literature to contrast the dual-control BCI paradigm with other paradigms.

To determine baseline tuning, variability and inferred connectivity of the recorded population, the monkey first performed a manual control (MC) target-pursuit task that involved acquiring targets through two degree-of-freedom wrist motion. When compared to this manual control task, we observe widespread changes in tuning of both control and non-control units for the brain-control and dual-control tasks. We find that in the dual control task both tuning to wrist movement and measures of effective connectivity between co-tuned units are more similar to those observed during natural motion. A notable exception is that effective connections to control units significantly decrease during the dual-control task. This suggests that the dual-control task requires control unit-specific dissociation from other units. We further find that the intrinsic variability of the control units has a significant effect on how the task is performed. Specifically, at least one control unit must have a high intrinsic variability to allow cursor control and target acquisition. The degree of co-tuning and strong effective connectivity with other units has no effect on the dual control task proficiency.

Results

Dual control task better maintains tuning to wrist motion compared to brain control task

In our previous study [22], we observed that a monkey was able to use all tested pairs of single cortical units to control a BCI task, independently of ongoing contralateral limb movement. Overall performance in this dual-control task was similar, regardless of the control units' directional tuning. During both the brain control and dual control tasks, we measure how units change their tuning to wrist motion and compare this tuning to that observed during the manual control task. Our goal was to understand the overall differences in neural activity that underlie performance of the brain-control task compared to the dual-control task where wrist motion is also required. We characterize the relationship between unit activity and wrist motion by fitting a simple linear model which provides a measure of a unit's correlation with wrist velocity, given by a tuning angle and tuning strength (see Methods). We then analyzed changes in tuning to wrist motion between different tasks. Compared to the manual control task, we observed significant changes in average tuning angle ($p < 0.001$; two sided t-test, Fig 2A,B) and average tuning strength ($p < 0.001$, Fig 2C) for units in both the brain control and dual control task. To test the possibility that 90 degrees of visual feedback rotation could by itself result in observed changes in tuning angle, we also conducted a set of trials in which visual feedback was not rotated (Fig S2). In both conditions we observed similar results, suggesting that tuning changes are due to the dual control task rather than the visual rotation.

We asked whether control unit tuning was specifically modified in the two tasks. We found that during the brain control task, 56% of control units and 72% of non-control units did not significantly change their tuning angle compared to the manual control task. However, in dual control, unit tuning angles were more constrained by wrist motion: 79% of control units and 81% of the non-control units retained a similar tuning angle as in manual control (Fig 2B). Thus, the significantly greater change in the tuning angle of the control units compared to non-control units

in the brain control task (p=0.008, two sided t-test) was not observed in dual control, where tuning angle changed similarly for both types of units (p=0.699, two sided t-test).

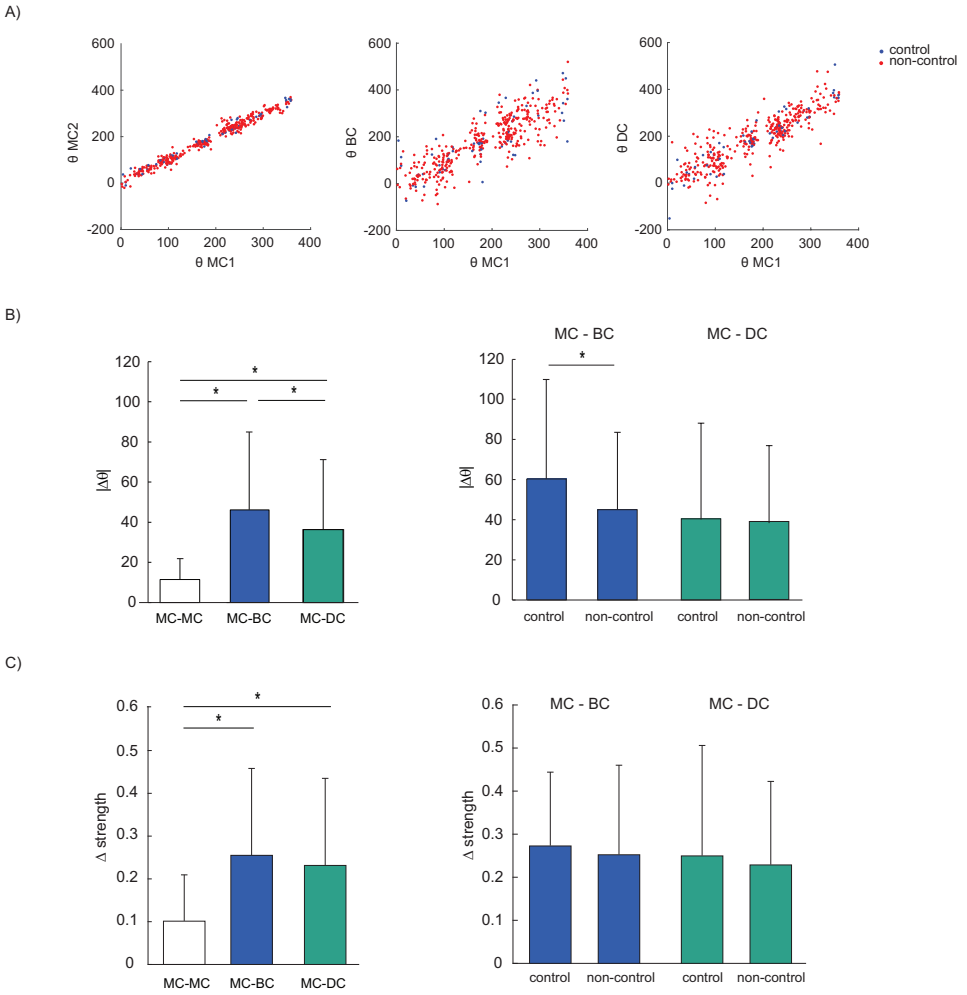

Fig 2. Changes in encoding between dual control and brain control tasks. A: Change in each unit's tuning angle (θ) between conditions. Each point represents one unit within one set of recordings. All angles on the y axis shifted to within ± 180 degrees of x axis value. Plot shows 69 days of recordings, with a total of 597 units in all sessions and an average of 8.652 units per session. Only units tuned, with a R^2 above 0.01, were included. Compared to the manual control task (MC1), the population of units showed significant changes in their tuning angle in both the brain (BC) and dual control task (DC). Unit tuning angle is more constrained by wrist motion in dual control than in brain control task. During the second manual control session (MC2), the majority of units reverted back to their original tuning angle in MC1. B: (left) Change in tuning angle (θ) for all units between conditions. Units changed tuning angle significantly between brain control and manual control and between dual control and manual control task (Two sided t-test; both $p \ll 0.001$). (right) Change in tuning angle for control and non-control units between conditions. Units controlling the BCI change their tuning angle significantly more than non-control units during BC (t-test; $p = 0.008$). In the dual control task, BCI units did not change their tuning angle significantly more than non-BCI units ($p = 0.699$). All values are mean ± standard deviation. C: (left) Change in tuning strength for all units between conditions. Compared to the manual control task (MC1), the population of units showed significant changes in their tuning strength in both the BC and DC task ($p \ll 0.001$). (right) Tuning strength changes similarly for control and non-control units in both tasks.

To assess whether the tuning changes in previous brain control trials persisted, the monkey performed another manual control task after the dual control task. During this post-manual control session, the majority of units had the same tuning angle as in the first manual control task (correlation coefficient 0.99; Fig 2A), suggesting that the network dynamically reconfigured in a reversible way. This reversibility is similar to that seen in previous BCI studies [20].

Dual control task selectively dissociates control unit activity from co-tuned units

We next considered how the relationships between pairs of units depended on the task. To evaluate this, we used a measure of effective connectivity, transfer entropy. The transfer entropy of two random processes, from X_t to Y_t, measures the information about Y_t obtained by observing the history of X_t, conditioned on the history of Y_t (see Methods). In this paper we use the term effective connectivity to mean a directed relationship in which one neuron's spike train is useful in predicting another's [37,38]. We note that, though directed, transfer entropy is purely a statistical relationship and does not imply any synaptic or anatomical connection.

We investigated the relation between control units and co-tuned unit activity during the dual control task. As described above, in dual control most units maintain their association with wrist motion observed in manual control. However, the control units are required to modulate their activity independently of this association. Thus we might expect that they will subsequently exhibit a weaker relation with units with which they were previously co-tuned, and that this will manifest as a change in the transfer entropy between control and co-tuned units. According to this hypothesis, the dual-control task requires control unit activity to become specifically less dependent on co-tuned neural activity. We therefore used a directed measure of connectivity to measure this dependence.

We focused on changes in connectivity that occurred between manual control and brain or dual control tasks for different subpopulations of unit pairs (see Methods). Overall, we observed a widespread change in the transfer entropy between manual control and the brain and dual control conditions. This change was smallest between the pre- and post- manual conditions ($R^2 = 0.420$, with the brain and dual control conditions showing larger changes in transfer entropy (Fig 3A). The changes for both brain control ($R^2 = 0.212$) and dual control ($R^2 = 0.221$) were similar (Fig 3A).

We then compared the change in connectivity for control units versus non-control units. We observed that between brain and manual control, control units decreased their connectivity significantly more than non-control units. We observed the same trend when comparing changes in connectivity between dual and manual control (Fig 3B). Thus, in both brain and dual control tasks, the control units become more independent from other units than non-control units. In other words, both tasks result in some control unit-specific dissociation.

We wondered if this control unit dissociation was specific to units that were co-tuned with the control unit. Indeed, comparing changes in connectivity to the control unit between brain and manual control showed a significant difference in the connectivity change for previously co-tuned units versus randomly selected units ($p \ll 0.001$ [rank sum]; Fig 3C). However, comparing such changes in connectivity to the control unit between dual control and manual control does not reveal a significant difference ($p = 0.529$ [rank sum]). Thus, control units become more independent from co-tuned units than from other units in the brain control task, but not in the dual control task.

This result is counterintuitive, given that the structure of the dual-control task was designed to effect the opposite result – control unit specific independence. The possibility remains, however, that the brain control task leads to a higher degree of change in co-tuned units in general, and that the effect has nothing to do with whether or not a control unit is involved. Thus, we also examined changes in connectivity to non-control units, from either co-tuned units or from non-co-tuned units (Fig 3D). In the brain control task, connectivity decreases significantly more for co-tuned units than for non-co-tuned units. This is similar to the result obtained in Fig 3C for control units. In dual control, however, connectivity changes significantly less for co-tuned units than for non-co-tuned units. In fact, between manual and dual control the overall change in connectivity for co-tuned units to non-control units is not significantly different from zero ($p = 0.145$). Thus, in dual control

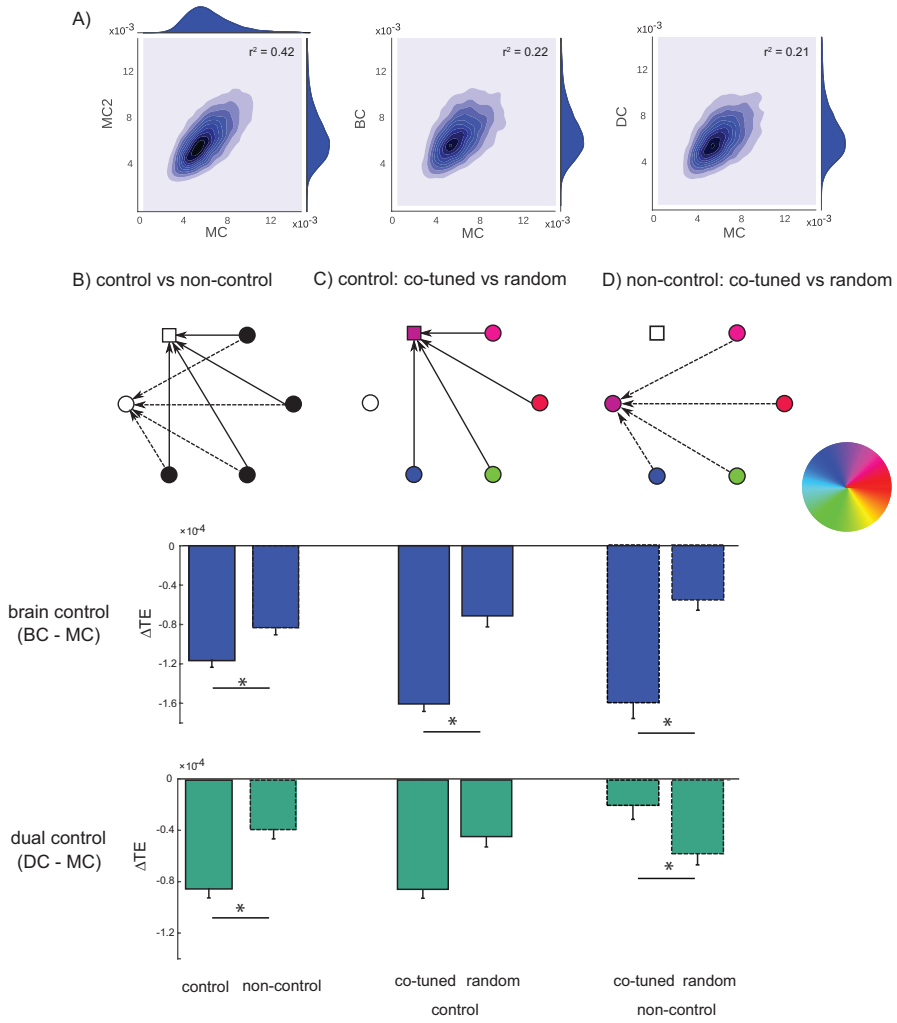

Fig 3. Connectivity analysis using transfer entropy. A: Contour density plot showing change in transfer entropy between tasks. Changes between first and second manual control tasks $R^2 = 0.423$, manual control to brain control $R^2 = 0.212$, manual control to dual control $R^2 = 0.221$. B: Top: Cartoon demonstrating comparisons in change in transfer entropy between tasks for an example, six unit network with one control unit. Square indicates control unit, circles indicate non-control units. Here effective connectivity to control (solid arrows) and to non-control units (dashed arrowed) are compared. Bottom: connectivity changes significantly more between manual and brain control task for control units than non-control units ($p \ll 0.001$ [Wilcoxon rank sum]); connectivity also changes significantly more between manual and dual control task for control vs non-control units ($p \ll 0.001$ [rank sum]). All changes are significantly different from zero (all $p \ll 0.001$ [signed rank test]) C: Change in transfer entropy to control units between tasks for both co-tuned units and randomly selected units. Here color indicates tuning angle, per the color wheel to right. Connectivity to the control units changed significantly more between manual and brain control task for co-tuned units than for randomly selected units ($p \ll 0.001$ [rank sum]); connectivity to the control units does not change significantly more between manual and dual control task for co-tuned units vs randomly selected units ($p = 0.528$ [rank sum]). All changes were significantly different from zero ($p \ll 0.001$ [signed rank test]) D: Change in transfer entropy to non-control units between tasks for both co-tuned units and randomly selected units. Connectivity changed significantly more between manual and brain control task for co-tuned units than for randomly selected units ($p \ll 0.001$ [rank sum]); connectivity changes significantly between manual and dual control task for randomly selected units compared with co-tuned units ($p \ll 0.001$ [rank sum]). Changes for co-tuned units between manual and dual control were not significantly different from zero ($p = 0.145$ [signed rank test]).

transfer entropy between co-tuned units is maintained, unless one of the units is a control unit. In brain control, however, there is an overall decrease in connectivity for co-tuned units, independently of whether a control unit is involved.

This result is consistent with the tuning angle findings of the previous section, in which greater changes in tuning angle were observed in brain control compared to the changes in dual control (Fig 2B). Together these findings suggest that the manual component of the dual-control task holds both the tuning angles and connectivities between co-tuned units in similar states to that observed in manual control, with the notable exception of connectivities involving the control units.

Neural variability constrains brain control proficiency

Despite the insight that the connectivity and tuning analysis provide into the neural activity that underlies the dual control task, neither analysis was predictive of task proficiency. This is consistent with previous studies showing remarkable flexibility of motor cortical activity. Previous studies also demonstrate that neural variability measured during natural behavior – intrinsic variability – is a constraint on learning and proficiency of brain control tasks [6]. Thus we investigated the role of intrinsic variability of the control neurons in relation to the variability of the entire neural population.

We began by identifying a low-dimensional space that characterizes a high proportion of the variance observed during natural movements from the manual control recordings. We used Gaussian process factor analysis (GPFA) to identify this space, which we term the intrinsic space (similar to the intrinsic manifold proposed by Sadtler et al 2014 [6]). In line with this terminology, we term the variance of each recorded unit during the manual control task its intrinsic variability. When a brain control task is performed, the activity from just two neurons determines the cursor position. The activity of these two neurons thus comprises a brain-control space (Fig 4A). Based on Sadtler et al 2014 [6], we hypothesized that more proficient brain-control and dual-control would occur when the brain-control space and the intrinsic space share some amount of alignment. We observed that in the brain control task, but not in the dual control task, performance is significantly higher when the brain-control and intrinsic spaces are indeed non-orthogonal (Fig 4B; brain-control $p = 0.023$, dual-control $p = 0.13$ – two sided t-test).

Intrinsic variability of control units affects contributions to cursor control and task proficiency

Given this result, we sought further insight into how the task was being performed by focusing specifically on the intrinsic variability of the control units. In both the brain control and dual control tasks, the decoder (Eq 1) is set up so that neurons which have a high variance exert more effect on the cursor trajectory. We hypothesized that a high amount of variance is necessary for control over the cursor. If this is indeed the case then, when control units are selected with lower intrinsic variability, the monkey is required to increase unit variance in order to accurately perform the target pursuit task. If the monkey is unable to change the control units' variance, then we might expect that in these cases he will be less able to control the cursor and will acquire targets at a lower rate.

To measure each unit's contribution to the target acquisition task, we constructed a metric which quantifies the amount of influence each control unit has over the cursor trajectory. Since the cursor trajectory was determined through a moving average model weighing the contribution from each control unit, we selected a Granger causality-type metric (see Methods). We validated this metric using synthetic cursor data generated through a weighted contribution of two neurons (Figure S4). We observed that a high value of this control metric for at least one of the control neurons is required for dual control task proficiency (Fig 5A).

In addition, we observed that a high control metric is related to a high intrinsic variability of the control unit (Fig 5B). This indicates that the task is predominantly controlled by the unit with higher intrinsic variation. As suggested by the relation to cursor control and performance, there is

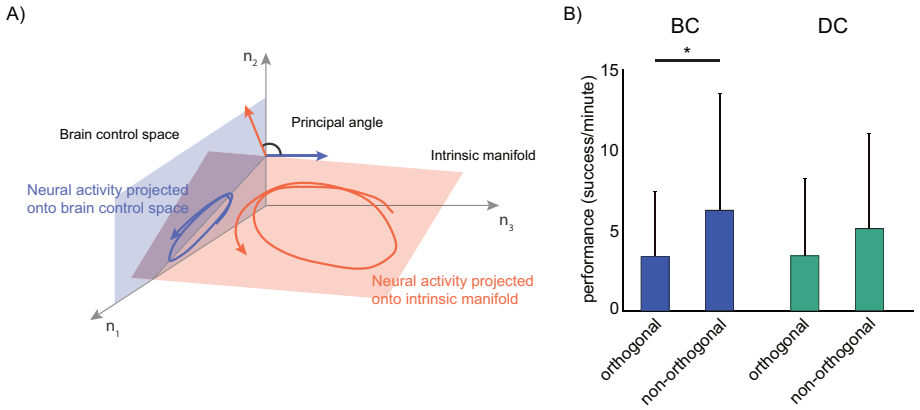

Fig 4. Intrinsic variability constrains performance. A: Cartoon demonstrating variability analysis. For simplicity here we assume only three neurons are recorded. At each recorded time bin, we collect three observations, the number of spikes occurring within the time bin of each of the recorded neurons, and place this point within a three dimensional 'neural activity space'. In this way we collect a cloud of points representing the activity of all neurons throughout the entire recording. We identify axes which represent a lower-dimensional subspace which accounts for the largest proportion of the variance of the data. Here Gaussian Process Factor Analysis (GPFA) is used to identify this subspace, and is used to identify such a low-dimensional space in conjunction with an assumption of temporal continuity between successive recordings of each neuron. This space is called the 'intrinsic manifold', and here is constructed based on manual control recordings. Units chosen for brain control and dual control tasks provide an n-dimensional subspace, where n is the number of neurons used in the decoder (here $n = 2$). The alignment between the brain-control space and the intrinsic space is the 'principal angle'. For clarity, only the neural trajectories projected onto the two low-dimensional spaces are shown, not the full-dimensional neural trajectory. B: Alignment between brain-control space and intrinsic manifold relates to performance in brain control and dual control. In brain control, performance is significantly higher when the two spaces are non-orthogonal, i.e. when there is some degree of overlap between the two spaces. In dual control, the same trend exists, though it is not significant ($p = 0.13$; two sided t-test).

a significant, positive relation between control unit intrinsic variability and performance, in both brain control (not shown) and dual control (Fig 5C).

These results suggest that a control unit's intrinsic variability is a harder constraint on both brain-control and dual-control task performance than either its tuning or its transfer entropy with other units. Therefore, when selecting units to control a BCI, intrinsic variability should be given greater consideration than tuning or connectivity to other units.

Discussion

In this study we sought to understand the factors that affect dual-control BCI proficiency and that may inform control unit choice in concurrent BCI tasks. Together our analyses suggest that factors such as effective connectivity and control unit's tuning strength and direction have minimal impact on dual control performance, while the unit's intrinsic variability has a significant effect.

Differences between linear tuning in brain control and dual control

In both brain and dual control contexts we observed a population-wide change of tuning angles between conditions, compared to tuning to naturalistic motion. These changes are suggestive of

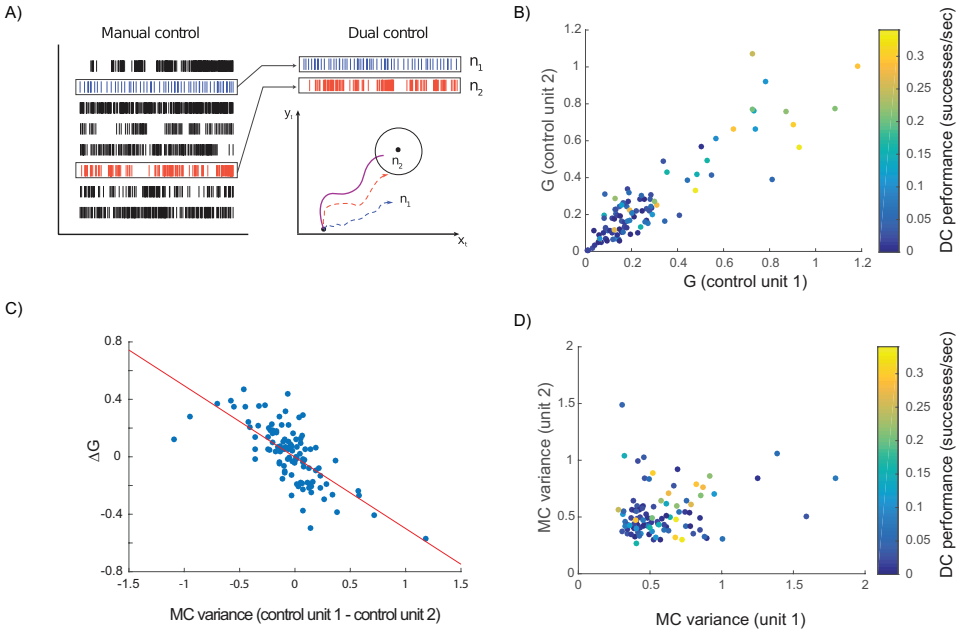

Fig 5. Contributions of cursor motion by each control unit, measured by Granger causality score $\mathcal{G}$. A: Control units chosen based on their tuning and activity in manual control recording (left panel) exert different levels of control over cursor in dual control recording (right panel). Dashed lines represent cursor motion contribution from individual control units, while solid purple line represents their aggregate. B: High levels of control positively correlated with high performance indicated by warmer colors clustering to the right (F-test, $p \ll 0.001$). C: Granger causality analysis identifies extent that individual control units exert control over cursor trajectory. The relative variance of the control units in manual control (MC) significantly influences the amount of control that each unit exerts over the cursor trajectory in dual control (DC), as identified by the Granger causality analysis ($p \ll 0.001$, $R^2 = 0.500$). D: The variance of control units observed during manual control is significantly positively correlated with performance in the dual control task. (F-test, $p = 0.016$).

a global reconfiguration of the motor cortical network for different motor tasks, as seen in many previous BCI studies [14, 16–21].

We also observed that tuning angles are significantly different between brain control and dual control. This finding is similar to observations made in a previous concurrent control study [23]. The authors found that producing static force with the arm during the BCI task perturbs the neural map formed during the standard BCI task. Interestingly, we found that during dual control, tuning angles were more similar to tuning observed during natural motion. Consistent with this finding, Suminski et al [39] reported that when the arm is passively moved in concert with a brain-controlled cursor, decoding performance improved. This suggests that tuning curves change less between manual and brain control conditions when the arm is actively involved in the task. It is possible that proprioceptive feedback inherent in arm movement is responsible for driving uniformly distributed tuning angles aligned with arm movement directions.

In contrast to the brain control task, control-unit-specific changes in tuning were not observed in the dual control task. Greater modification of control units' tuning angles in the brain control task appears to be consistent with results from other studies in which visuomotor rotation was applied to a subset of units controlling the cursor [18, 40]. Applied visuomotor rotation resulted in the rotated subpopulation shifting their tuning angles more than the non-rotated population. Conversely, in the dual control task, tuning angles changed similarly for both control and non-control units, as observed in [5]. One explanation for this difference is that the manual component of the dual

control task constrains neural activity in such a way that dissociation can only be achieved at a broader population level, or that the increased task complexity requires more coordinated activity of neurons in the population.

Previous BCI studies additionally show differential tuning strength of control versus non-control units [12,20], while here we report that both control and non-control units undergo similar changes in tuning strength in the brain control and dual control task. Interestingly, Ganguly et al. observed no differential tuning strength of control and non-control groups of units during the initial learning process, but only after 2-3 days of performance with the same units. Similarly, Law et al [12] saw a stronger effect in late compared to early trials. This suggests that the short amount of practice time in the present study for each pair of control units may be insufficient to affect modulation depth.

Insights from effective connectivity

We sought a better understanding of the changes in network activity underlying the dual control task compared with the brain control task. We observed a selective decrease in effective connectivity between co-tuned units and the control unit in dual control, and a non-selective decrease in connectivity between co-tuned units in brain control. This suggests that the dissociation required in the dual control task does affect interactions between individual units in the network.

Effective connectivity, sometimes also referred to as functional connectivity, has been used to study properties of motor cortical networks using directed information in [41, 42] and the Granger causality approach of [43]. In this study we chose to use transfer entropy as a non-parametric measure of effective connectivity, as it performed best on a synthetic validation dataset. A drawback of all these methods is that they are unable to distinguish between direct interactions and the effect of hidden, or latent inputs to recorded units. Given the 400 μm electrode separation of the Blackrock Utah array used here, direct synaptic connections between units are unlikely [44] and all relations between units may be mediated through unobserved connections. A more appropriate model then may be one that separates latent and direct connectivity components [45–47]. Such an analysis is the focus of future work.

Intrinsic variability constraints

In the brain control task we found that performance is constrained by each unit's variability, and how well aligned the brain control space is to the units' natural firing patterns. Other studies have also reported that factors relating to motor tuning may not be the most relevant when selecting BCI control units. For instance, Yu et al 2009 showed that BCI performance was better when low-dimensional latent dynamics were inferred and utilized, independent of any kind of external movement or task parameter [33]. Another study showed that trial-by-trial spike predictions based only on the firing rates of simultaneously recorded motor cortex neurons tend to outperform predictions based on external parameters [48].

Constraints provided by intrinsic neural variability have been identified before [5, 6] in populations of 10s to near 100 units. In this study the 'control space' is only two dimensional, whereas previous studies have explored the concept in wider control spaces consisting of 10-100 units. The low-dimensionality of the present study may facilitate insight into these larger control spaces. In Sadtler et al 2014 [6], for instance, within-manifold perturbations were less detrimental to performance than 'outside manifold' perturbations. In the present study the 'control space' is deliberately limited to two dimensions to facilitate insight into how a focused population of neurons must change their activity to accomplish the dual control task. It is possible that the role of variability of control units also explains these previous results. That is, the intrinsic manifolds are predominantly defined by axes parallel to highly varying/firing units, and it is these units that dominate the manifold, not linear combinations of units. In other words, correlations between units may not be significant; within/outside manifold perturbations become simply the inclusion or otherwise of varying units that are able to contribute to the control of the BCI. This possibility requires further study.

Outlook

Our results demonstrate that dissociation of motor cortical units required by a dual control task does not occur with individual neuron specificity, but through task-related changes in both the control and non-control populations. This and previous work [5, 6] suggests that, provided specific constraints are taken into account, the motor cortex can flexibly adapt to a variety of challenging control tasks. Tasks that require dissociations from established correlations between neural activity and movement can be learned even over the course of a short recording session. Alongside their clinical value, BCI paradigms provide insight into the physiological principles that guide motor control [25] by allowing the relevant populations for motor control and learning to be directly observed. Despite recent findings, a lack of consensus remains over the appropriate framework to understand primary motor cortical activity. In the present study we relied on a variety of analyses utilized in both representational and dynamical perspectives of encoding in motor cortex. We found that each analysis revealed different insights into how primary motor cortex learns a novel coordination task.

Materials and methods

One macaque nemestrina monkey was trained to perform a random target-pursuit motor task [22]. He began each daily session by controlling the cursor with isometric wrist torque (manual control, MC), then progressed to using the aggregate neural activity of two single units to control a cursor in one dimension (brain-control, BC). Subsequently, he used the same neural activity to control the BCI in one dimension, while simultaneously using isometric wrist torque of the contralateral forelimb to control the cursor in a second orthogonal dimension [22] (dual-control, DC; Fig 1). All procedures were in accordance with National Institutes of Health 'Guide for the Care and Use of Laboratory Animals'.

The monkey was required to move the cursor towards the target and maintain the cursor within the target for at least 1s to receive a reward. Targets appeared randomly, and a 0.5s break was provided between trials. For each trial there was a time-out period of 40s. Each day a new pair of units was chosen to control the BCI. If directionally tuned units were used that day, units with approximately opposite preferred directions were selected and paired to control the BCI.

In order to ensure decoupling of unit activity and hand control in the dual control task, visual feedback of one of these modalities had to be rotated to achieve independent degrees of freedom on the monitor. Since the monkey was overtrained on the manual control task, we made a deliberate decision to preserve the relation between wrist torque and cursor movement. Instead, we rotated the visual feedback of the units' firing rate by 90 degrees relative to their preferred direction, even for units with little to no directional preference.

All units with a mean firing rate above 5 Hz recorded from the hemisphere contralateral to arm movement were analyzed. The torque signals were smoothed with a Gaussian kernel of width 300ms and a linear tuning model was fit to each unit's firing rate as a function of wrist torque. Overall performance was quantified as the number of targets per minute acquired throughout the entire recording.

BCI decoder

The cursor position along the BCI axis, x_t, was determined as a linear function of N control neurons' smoothed activity, y_t^n, $N = 2$, with gain α_n and a running estimate of baseline firing b_n^t:

$$x_t = \sum_{n=1}^{N} \alpha_n (y_t^n - b_t^n) \tag{1}$$

where the baseline firing rate is defined as

$$b_t^n = b_{t-1}^n (1 - \gamma) + y_t^n \gamma \tag{2}$$

The baseline update rate is given by γ ($0 < \gamma < 1$). Spikes were binned at 60Hz. Spike trains were smoothed with a Gaussian kernel of width 100ms.

Linear tuning model

The kinematic encoding model to wrist motion is defined for each unit y^n simply through the linear relation

$$y_t^n = \alpha^1 x_{t+\tau}^1 + \alpha^2 x_{t+\tau}^2 + c + \epsilon \tag{3}$$

for velocity along task axes x^1 and x^2, for predetermined time lag τ, and for Gaussian noise ϵ. Here an offset of $\tau = 75$ms was used (Fig S5). Spikes are binned at 25 Hz. This time is consistent with previous studies [30] and was, on average, the time lag of the maximum cross-correlation between cursor and neural activity. For most units, a velocity encoding model was found to yield higher cross-correlation between torque and neural activity than a position encoding model, so the results presented here are based on a velocity encoding model. Velocity is computed from position information using a cubic spline [30].

Note that the model is equivalent to a cosine encoding model through a simple transformation:

$$y_t^n = \alpha r \cos(\theta - \theta_{pref}) + c + \epsilon \tag{4}$$

for $\alpha = \sqrt{(\alpha^1)^2 + (\alpha^2)^2}$, $\theta_{pref} = \tan^{-1}(\alpha^2/\alpha^1)$, $r = \sqrt{(x^1)^2 + (x^2)^2}$ and $\theta = \tan^{-1}(x^2/x^1)$. From this it is clear we can interpret α/c as a measure of modulation depth. We will refer to θ_{pref} as tuning angle and α as tuning strength.

Assessing effective connectivity with transfer entropy

Transfer entropy [31] is defined as the conditional mutual information between an observed time series Y_t and the history of a candidate related series X_t, conditional on the history of X_t:

$$H_{X\to Y} = I(Y_t|Y_{t-1}, Y_{t-T}) - I(Y_t|Y_{t-1}, Y_{t-T}, X_{t-1}, X_{t-T}) \tag{5}$$

As with many information theoretic quantities, $H_{X\to Y}$ is expensive to compute. An approximation tailored for spiking data is used, in particular the transfer entropy toolbox [32].

The method was validated using a synthetically generated dataset and compared with other common effective connectivity measures. Specifically it was compared to linear Granger causality, Poisson-process Granger causality, and correlation. On the synthetic data transfer entropy was most accurate. More details are provided in the supplementary methods.

Transfer entropy is computed with spikes binned at 5ms intervals. Six minutes of recording data in each condition is used. Connections extend up to 30 time-bins into the past (150ms).

The connectivity analysis is performed as follows. Within each session, between recording conditions (MC→BC and MC→DC), changes in connectivity are compared for different populations of pairs of units. Recalling transfer entropy is a directed measure of connectivity, the populations of pairs of units compared are: to control units and to non-control units; co-tuned units to control units and non-co-tuned units to control units; and co-tuned units to control units and co-tuned units to non-control units. For a given session, generally no more than two pairs of units could be identified within 45 degrees of the unit of interest (e.g. the control unit). Thus in comparisons involving co-tuned units, for each session only the top two co-tuned units were selected. These are defined as the two unit pairs having the closest tuning angle. All other pairs of units in the session count as not co-tuned. In non-co-tuned populations generally many more than two pairs of units can be identified within a session. Thus to ensure these larger populations are better sampled, but still balanced with populations of co-tuned units, the selection of two units is bootstrap sampled 50 times for each session. Such generated populations are combined over all sessions. All co-tuned populations are selected on the basis of tuning during manual control recordings.

Intrinsic variability and Gaussian-process factor analysis

Gaussian process factor analysis [33] was used to identify a 15 dimensional subspace in the manual control recordings, consistent with previous analyses into the dimensionality of motor cortical data [33, 34]. Data were binned at 2ms intervals. For each recording, 300s of data were used. Of the entire 15-dimensional space identified, the top five dimensions were shown to accommodate for, on average, 98% of neural covariability in the data. The principal angle was identified using the subspace function in MATLAB.

Granger-causality cursor control metric

Given the form of the linear decoding model (Eq 1), it is reasonable to model the cursor trajectory as a moving average time series:

$$x_t = \sum_{n=1}^{N} \sum_{p=1}^{P} k_p y_{t-p}^n + c \tag{6}$$

for some model order P to be determined. Spikes are binned at 25Hz, and cursor data is smoothed with a cubic spline, as per the linear tuning model [30]. From the linear model we compute the maximum likelihood estimate given both neurons, $\hat{\beta}_{MLE}^n$, and the maximum likelihood estimate not given neuron j, $\hat{\beta}_{MLE}^{n/j}$. The difference between the two log-likelihoods of the MLE models provides a Granger-causality type [35] metric

$$\mathcal{G}_{j\to i} = 2\mathcal{L}(\hat{\beta}_{MLE}^i; \mathbf{x}, \mathbf{y}) - 2\mathcal{L}(\hat{\beta}_{MLE}^{i/j}; \mathbf{x}, \mathbf{y}) \tag{7}$$

that quantifies the influence that one neuron has on the cursor trajectory. The likelihood ratio is often used as the basis for hypothesis testing, though the quantity itself can be interpreted as a form of transfer entropy, provided the data is Gaussian distributed [36]. The Granger causality toolbox MVGC [35] was used to perform this analysis. The approach is well suited since the cursor position is, by definition, a moving average model of the neural data. The method assumes a covariance stationary time series. By construction (Eq 1) the cursor trajectory has this property, provided that the neural data is second-order stationary. This was confirmed by checking the spectral radius of the vector autoregressive model was less than one.

To verify the notion that $\mathcal{G}_{j\to i}$ has quantitative meaning we generated synthetic data from neural recordings. A cursor trajectory was generated as the weighted sum of the two units' decoded trajectories, according to (Eq 1). Following this we computed $\mathcal{G}_{j\to i}$ for different weightings. A strong monotonic relationship exists between $\mathcal{G}_{j\to i}$ and the weight that determines how much each unit contributed to the cursor trajectories (Fig S4), demonstrating that the control metric can recover which unit contributes more to cursor control.

Analysis

Analysis was performed in MATLAB 2016a. Code is available online at https://github.com/benlansdell/dualbci.

Supporting information

S1 Appendix. Validating connectivity methods.

There are a number of candidate measures of effective or functional connectivity between units, including Granger causality, transfer entropy and correlation. Undirected, statistical relationships, may be referred to instead as functional connectivity [37, 38]. Before applying these analyses to brain-control recording data we sought to understand each method's success in identifying effective/functional connectivity in a synthetic dataset with known connectivity. For this, we generated a dataset using a point process model [49]. This model was designed to simulate an ensemble of cortical neurons in motor cortex, and has been used previously for making predictions about brain-BCI

interactions. The candidate methods were standard covariance, Granger causality analysis based on a Poisson process GLM [43], Granger causality analysis based on a linear firing rate model, and transfer entropy [32]. Details for each method are below.

A simulated spiking network [49] of 30 Izhikevich neurons is divided into 3 densely interconnected groups with some inter-group connections. As common input, the neurons are stimulated with a rectified sinusoidal input. The input was a 2 Hz sinusoidal envelope with a 100 Hz envelope amplitude. The simulation was run for 3 min, with a baseline firing rate of 1 Hz (producing on average 17 Hz activity in simulation).

For our Transfer entropy analysis, we used the Transfer Entropy Toolbox [32] with a 1ms bin size, a max delay of 20 ms, 1-3 time bins and 1-3 past time bins to calculate the transfer entropy between neurons. We then normalized the Transfer entropy by z-scoring.

For our linear Granger causality analysis, using MLE we fit a linear model of the activity of all neurons 1-20 ms in the past (including the past history of the unit being predicted) in order to predict the activity of unit i. We then removed unit j from the model and re-fit the model. The change in deviance (normalized by z-scoring) between the old and new model was used as our measure of the effective connectivity between units.

For our GLM granger causality analysis, we performed the same analysis with an Poisson model with an exponential nonlinearity rather than a linear model, again normalizing our results by z-scoring.

We also calculated the covariance between units and normalized by z-scoring.

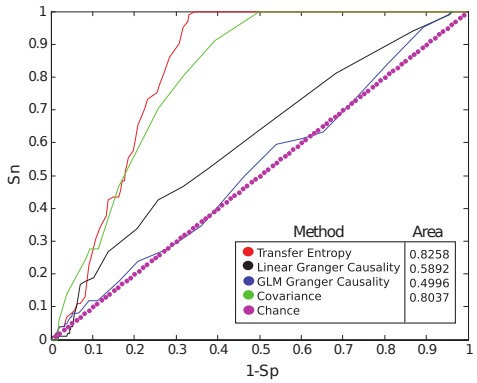

Fig S1. Validation of effective/functional connectivity models. ROC curve comparing performance of functional/effective connectivity inference methods on simulated spike trains. The transfer entropy has the largest ROC area for in simulated data. Refer to S1 Appendix for more information.

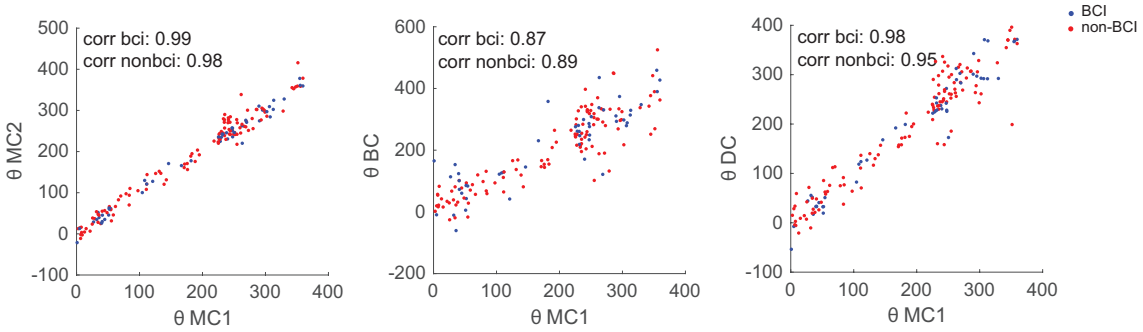

Fig S2. Change in tuning angle in unrotated sessions. Change in tuning angle between conditions for control units that were displayed in their preferred direction on the screen, rather than rotating by 90 degrees as in the remainder of the study to permit dual control. Changes are similar to those observed in the 90 degree rotated dataset (Fig 2) – overall larger between manual control and brain control tasks than manual control and dual control tasks.

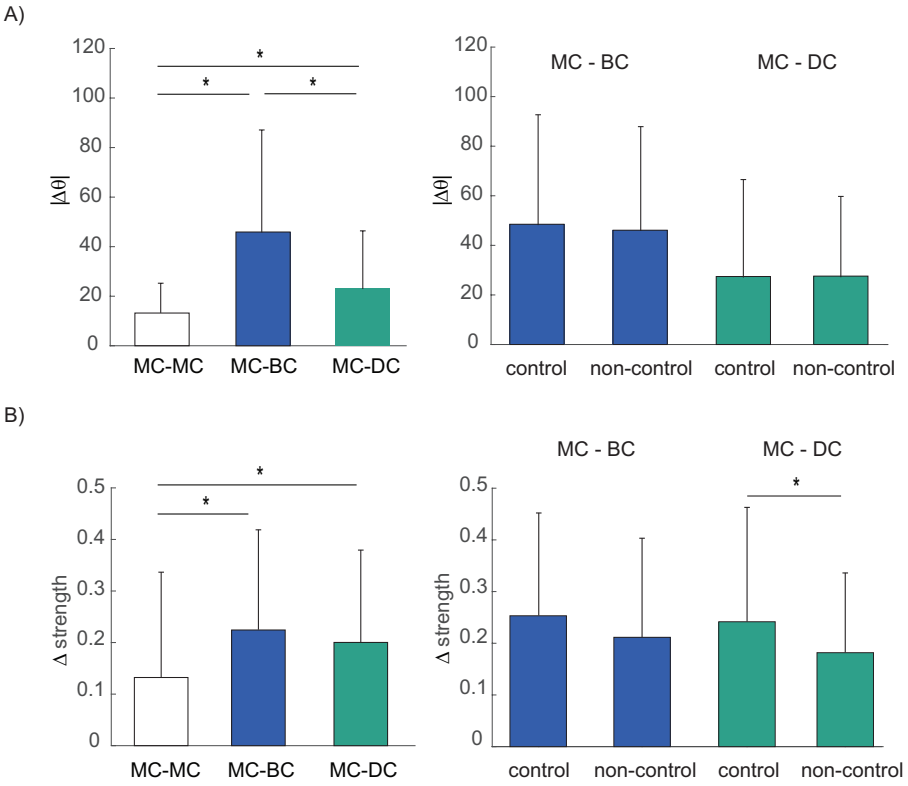

Fig S3. Tuning changes during unrotated sessions. A: Tuning angle. Left: all p $\ll 0.001$ [two sided t-tests]. Right: non-significant. B: Tuning strength. Left: MC-BC $p \ll 0.001$, MC-DC $p = 0.0016$ [two sided t-tests] Right: MC-DC $p = 0.0491$ [two sided t-test]

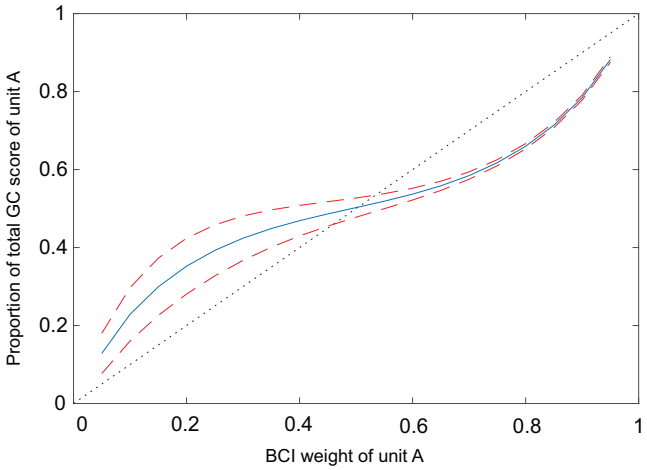

Fig S4. Validation of Granger-causality cursor control metric. Synthetic cursor trajectories are generated according to neural data from two control units, according to the decoder (1). Contribution weights for each units are varied. 108 trajectories are generated for each weighting and cursor control metric, $\mathcal{G}_{j\to i}$, is computed. Blue curve indicates mean proportion of control score as a function of weight; red curve indicates standard deviation; dotted black line indicates identity line.

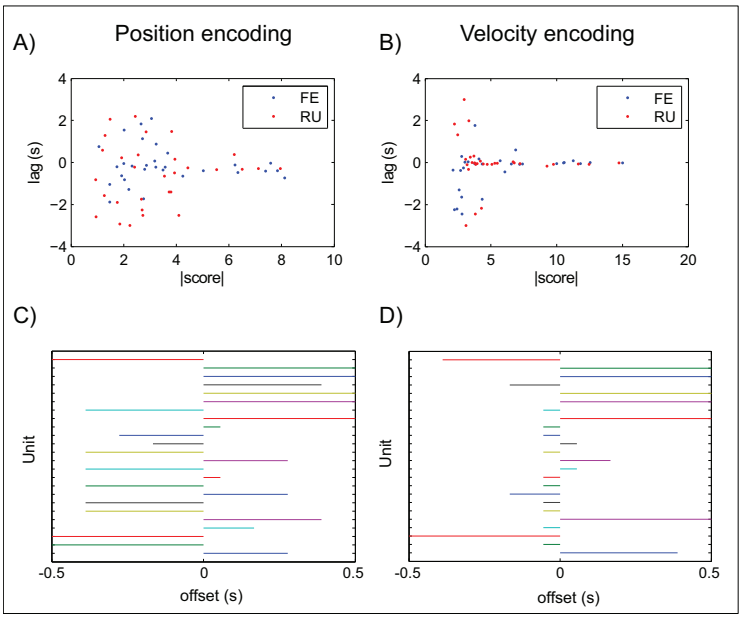

Fig S5. Optimal time lags for simple tuning models. Peak of cross-correlation between filtered firing and position flexion-extension (FE) and radial-ulnar (RU) torque output gives an optimal encoding lag for each channel. The absolute value of the z-scored peak provides a simple encoding 'strength'. A: The relation between score and optimal lag for direct FE and RU output – related to cursor positions. B: the same relation for time-differenced torque output – related to cursor velocity. Scores below 4 are wildly varying in optimal lag, while scores above 4 are more tightly constrained to near zero. The 'velocity encoding' responses is optimal near a lag of zero. Similarly, fitting a simple GLM Poisson model with a stimulus filter of length 1 (only one time point) for different time point values queries at what time delay the neuron's firing is most informative about the stimulus. The change in each model's likelihood when the stimulus filter is included gives an indication of the encoding at this time lag. The optimal time lag can be computed by taking the maximum in this change in likelihood. C: The optimal time lag for a set of units under a position encoding model. D: Optimal time lag for a set of units under a velocity encoding model. The most common optimal time lag was approximately -100ms – indicating spiking occurs before the motor output. As a compromise between these two sets of analyses, a fixed time lag of -75ms was used in the main text.

Acknowledgments

The authors thank Paul House for expert array implantations, Rob Robinson for animal handing and data collection, Charlie Matlack for task design, and Larry Shupe for technical support.

This work was supported by an American Heart & Stroke Association Scientist Development Grant [NCRP 09SDG2230091]; The Paul G. Allen Family Foundation (Allen Distinguished Investigator Award); National Institutes of Health [NS12542, RR00166]; the Center for Sensorimotor Neural Engineering (CNSE), a National Science Foundation Engineering Research Center [EEC-1028725].

References

1. Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, et al. High-performance neuroprosthetic control by an individual with tetraplegia. The Lancet. 2013;381(9866):557–564. doi:10.1016/S0140-6736(12)61816-9.

2. Gilja V, Pandarinath C, Blabe CH, Nuyujukian P, Simeral JD, Sarma Aa, et al. Clinical translation of a high-performance neural prosthesis. Nature medicine. 2015;21(10):6–8. doi:10.1038/nm.3953.

3. Aflalo T, Kellis S, Klaes C, Lee B, Shi Y, Pejsa K, et al. Decoding motor imagery from the posterior parietal cortex of a tetraplecig human. Science. 2015;348(6237):906–910. doi:10.7910/DVN/GJDUTV.

4. Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, Vogel J, et al. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. Nature. 2012;485(7398):372–375. doi:10.1038/nature11076.

5. Hwang EJ, Bailey PM, Andersen Ra. Volitional control of neural activity relies on the natural motor repertoire. Current Biology. 2013;23(5):353–361. doi:10.1016/j.cub.2013.01.027.

6. Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, et al. Neural constraints on learning. Nature. 2014;512(7515):423–426. doi:10.1038/nature13665.

7. Fetz EE. Operant conditioning of cortical unit activity. Science (New York, NY). 1969;163(3870):955–8.

8. Fetz EE, Baker Ma. Operantly conditioned patterns on precentral unit activity and correlated responses in adjacent cells and contralateral muscles. Journal of neurophysiology. 1973;36(2):179–204.

9. Fetz EE. Volitional control of neural activity: implications for brain–computer interfaces. The Journal of Physiology. 2007;579(3):571–579. doi:10.1113/jphysiol.2006.127142.

10. Moritz CT, Perlmutter SI, Fetz EE. Direct control of paralysed muscles by cortical neurons. Nature. 2008;456(7222):639–642. doi:10.1038/nature07418.

11. Moritz CT, Fetz EE. Volitional control of single cortical neurons in a brain-machine interface. Journal of neural engineering. 2011;8(2):25017. doi:10.1088/1741-2560/8/2/025017.

12. Law AJ, Rivlis G, Schieber MH. Rapid acquisition of novel interface control by small ensembles of arbitrarily selected primary motor cortex neurons. Journal of neurophysiology. 2014;(585):1528–1548. doi:10.1152/jn.00373.2013.

13. Fetz EE, Finocchio DV. Correlations between activity of motor cortex cells and arm muscles during operantly conditioned response patterns. Experimental Brain Research. 1975;23(3):217–240. doi:10.1007/BF00239736.

14. Carmena JM, Lebedev Ma, Crist RE, O'Doherty JE, Santucci DM, Dimitrov DF, et al. Learning to Control a Brain–Machine Interface for Reaching and Grasping by Primates. PLoS Biology. 2003;1(2):e2. doi:10.1371/journal.pbio.0000042.

15. Sharma M, Gaona C, Roland J, Anderson N, Freudenberg Z, Leuthardt EC. Ipsilateral directional encoding of joystick movements in human cortex. In: 2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE; 2009. p. 5502–5505. Available from: `http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5334559`.

16. Taylor DM, Others. Direct cortical control of 3D neuroprostheticdevices. Science. 2002;296(5574):1829–1832. doi:10.1126/science.1070291.

17. Lebedev MA, Carmena JM, O'Doherty JE, Zacksenhouse M, Henriquez CS, Principe JC, et al. Cortical Ensemble Adaptation to Represent Velocity of an Artificial Actuator Controlled by a Brain-Machine Interface. Journal of Neuroscience. 2005;25(19):4681–4693. doi:10.1523/JNEUROSCI.4088-04.2005.

18. Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB. Functional network reorganization during learning in a brain-computer interface paradigm. Proceedings of the National Academy of Sciences of the United States of America. 2008;105(49):19486–19491. doi:10.1073/pnas.0808113105.

19. Ganguly K, Carmena JM. Emergence of a Stable Cortical Map for Neuroprosthetic Control. PLoS Biology. 2009;7(7):e1000153. doi:10.1371/journal.pbio.1000153.

20. Ganguly K, Dimitrov DF, Wallis JD, Carmena JM. Reversible large-scale modification of cortical networks during neuroprosthetic control. Nature Neuroscience. 2011;14(5):662–667. doi:10.1038/nn.2797.

21. Golub MD, Yu BM, Chase SM. Internal models engaged by brain-computer interface control. Conference proceedings : Annual International Conference of the IEEE Engineering in Medicine and Biology Society IEEE Engineering in Medicine and Biology Society Conference. 2012;2012:1327–1330. doi:10.1109/EMBC.2012.6346182.

22. Milovanovic I, Robinson R, Fetz EE, Moritz CT. Simultaneous and independent control of a brain-computer interface and contralateral limb movement. Brain computer interfaces (Abingdon, England);2(4):174–185. doi:10.1080/2326263X.2015.1080961.

23. Orsborn AL, Moorman HG, Overduin SA, Shanechi MM, Dimitrov DF, Carmena JM. Closed-Loop Decoder Adaptation Shapes Neural Plasticity for Skillful Neuroprosthetic Control. Neuron. 2014;82(6):1380–1393. doi:10.1016/j.neuron.2014.04.048.

24. Bashford L, Wu J, Sarma D, Collins K, Ojemann J, Mehring C. Natural movement with concurrent brain-computer interface control induces persistent dissociation of neural activity. In: BCI Society Meeting. CA, USA; 2016.

25. Golub MD, Chase SM, Batista AP, Yu BM. Brain–computer interfaces for dissecting cognitive processes underlying sensorimotor control. Current Opinion in Neurobiology. 2016;37:53–58. doi:10.1016/j.conb.2015.12.005.

26. Michaels JA, Dann B, Scherberger H. Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. PLoS Computational Biology. 2016;12(11):1–22. doi:10.1371/journal.pcbi.1005175.

27. Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP. Reorganization between preparatory and movement population responses in motor cortex. Nature Communications. 2016; p. 13239. doi:10.1038/ncomms13239.

28. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, et al. Neural population dynamics during reaching. Nature. 2012;487(7405):51–6. doi:10.1038/nature11129.

29. Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? Neuron. 2010;68(3):387–400. doi:10.1016/j.neuron.2010.09.015.

30. Shoham S, Paninski LM, Fellows MR, Hatsopoulos NG, Donoghue JP, Normann RA. Statistical Encoding Model for a Primary Motor Cortical Brain-Machine Interface. IEEE Transactions on Biomedical Engineering. 2005;52(7):1312–1322. doi:10.1109/TBME.2005.847542.

31. Schreiber T. Measuring information transfer. Physical review letters. 2000;85(2):461–464. doi:10.1103/PhysRevLett.85.461.

32. Ito S, Hansen ME, Heiland R, Lumsdaine A, Litke AM, Beggs JM, et al. Extending Transfer Entropy Improves Identification of Effective Connectivity in a Spiking Cortical Network Model. PLoS ONE. 2011;6(11):e27431. doi:10.1371/journal.pone.0027431.

33. Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. Journal of Neurophysiology. 2009;102(1):614. doi:10.1152/jn.90941.2008.

34. Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. Cortical activity in the null space: permitting preparation without movement. Nature Neuroscience. 2014;17(3):440–448. doi:10.1038/nn.3643.

35. Barnett L, Seth AK. The MVGC multivariate Granger causality toolbox: A new approach to Granger-causal inference. Journal of Neuroscience Methods. 2014;223:50–68. doi:10.1016/j.jneumeth.2013.10.018.

36. Barnett L, Barrett AB, Seth AK. Granger Causality and Transfer Entropy Are Equivalent for Gaussian Variables. Physical Review Letters. 2009;103(23):238701. doi:10.1103/PhysRevLett.103.238701.

37. Nigam S, Shimono M, Ito S, Yeh FC, Timme N, Myroshnychenko M, et al. Rich-Club Organization in Effective Connectivity among Cortical Neurons. Journal of Neuroscience. 2016;36(3):670–684.

38. Sporns O. Brain connectivity. Scholarpedia. 2007;2(10):4695.

39. Suminski AJ, Tkach DC, Fagg AH, Hatsopoulos NG. Incorporating Feedback from Multiple Sensory Modalities Enhances Brain-Machine Interface Control. Journal of Neuroscience. 2010;30(50):16777–16787. doi:10.1523/JNEUROSCI.3967-10.2010.

40. Chase SM, Kass RE, Schwartz AB. Behavioral and neural correlates of visuomotor adaptation observed through a brain-computer interface in primary motor cortex. Journal of neurophysiology. 2012;108(2):624–44. doi:10.1152/jn.00371.2011.

41. Quinn CJ, Coleman TP, Kiyavash N, Hatsopoulos NG. Estimating the directed information to infer causal relationships in ensemble neural spike train recordings. Journal of Computational Neuroscience. 2011;30(1):17–44. doi:10.1007/s10827-010-0247-2.

42. Zhe Chen Z, Putrino DF, Ba DE, Ghosh S, Barbieri R, Brown EN. A regularized point process generalized linear model for assessing the functional connectivity in the cat motor cortex. In: 2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society. vol. 2009. IEEE; 2009. p. 5006–5009. Available from: http://www.ncbi.nlm.nih.gov/pubmed/19965032http://www.pubmedcentral.nih.gov/articlerender.fcgi

43. Kim S, Putrino D, Ghosh S, Brown EN. A Granger Causality Measure for Point Process Models of Ensemble Neural Spiking Activity. PLoS Computational Biology. 2011;7(3):e1001110. doi:10.1371/journal.pcbi.1001110.

44. Smith WS, Fetz EE. Synaptic Linkages Between Corticomotoneuronal Cells Affecting Forelimb Muscles in Behaving Primates. Journal of Neurophysiology. 2009;102(2).

45. Pfau D, Pnevmatikakis EA, Paninski L. Robust learning of low-dimensional dynamics from large neural ensembles; 2013.

46. Semedo J, Zandvakili A, Kohn A, Machens CK, Yu BM. Extracting Latent Structure From Multiple Interacting Neural Populations; 2014.

47. Macke JH, Buesing L, Cunningham JP, Yu BM, Shenoy KV, Sahani M. Empirical models of spiking in neural populations; 2011.

48. Stevenson IH, Kording KP. How advances in neural recording affect data analysis. Nature Neuroscience. 2011;14(2):139–142. doi:10.1038/nn.2731.

49. Lajoie G, Krouchev N, Kalaska JF, Fairhall A, Fetz EE. Correlation-based model of artificially induced plasticity in motor cortex by a bilateral Brain-Machine Interface. PLoS Comput Biol. 2017;13(2):1–35. doi:10.1371/journal.pcbi.1005343.